


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
 The ACM Digital Library The Guide

[Home](#) [About the Portal](#) [Search](#) [Advanced Search](#)
[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Published before June 2000

Terms used [profiling](#) [user](#) [advertisement](#)

Found 104 of 104,557

Sort results by relevance [Save results to a Binder](#)Try an [Advanced Search](#)Display results expanded form [Search Tips](#)Try this search in [The ACM Guide](#) Open results in a new window

Results 1 - 20 of 104

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [next](#)

Relevance scale

1 User profiling in personalization applications through rule discovery and validation

Gediminas Adomavicius, Alexander Tuzhilin

August 1999 **Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining**Full text available: [pdf\(674.55 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

2 Representation of electronic mail filtering profiles: a user study

Michael J. Pazzani

January 2000 **Proceedings of the 5th international conference on Intelligent user interfaces**Full text available: [pdf\(513.83 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Electronic mail offers the promise of rapid communication of essential information.

However, electronic mail is also used to send unwanted messages. A variety of approaches can learn a profile of a user's interests for filtering mail. Here, we report on a usability study that investigates what types of profiles people would be willing to use to filter mail.

Keywords: mail filtering, user studies

3 Delivery of electronic news: a broadband application

F. J. Burkowski, Michael A. Shepherd, C. R. Watters

October 1994 **Proceedings of the 1994 conference of the Centre for Advanced Studies on Collaborative research**Full text available: [pdf\(25.31 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes an ongoing research program for the development of an electronic news delivery system that exploits the promised high-bandwidth, switched, interactive communication facilities of the information highway. The research program is in a preliminary stage and will investigate the design and organization of the news sources, client/server architecture, and user interfaces leading to a prototype model electronic news delivery system. Initially based on a newspaper metaphor, the sys ...

4

Agent-mediated electronic commerce: issues, challenges and some viewpoints



h

c

g

e

cf

e

f

g

e

e

e

e

Hyacinth S. Nwana, Jeff Rosenschein, Tuomas Sandholm, Carles Sierra, Pattie Maes, Rob Guttmann
May 1998 Proceedings of the second international conference on Autonomous agents
Full text available: [pdf\(1.13 MB\)](#) Additional Information: [full citation](#), [citations](#), [index terms](#)

5 **InfoSleuth: agent-based semantic integration of information in open and dynamic environments**

R. J. Bayardo, W. Bohrer, R. Brice, A. Cichocki, J. Fowler, A. Helal, V. Kashyap, T. Ksiezyk, G. Martin, M. Nodine, M. Rashid, M. Rusinkiewicz, R. Shea, C. Unnikrishnan, A. Unruh, D. Woelk
June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2

Full text available: [pdf\(1.69 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The goal of the InfoSleuth project at MCC is to exploit and synthesize new technologies into a unified system that retrieves and processes information in an ever-changing network of information sources. InfoSleuth has its roots in the Carnot project at MCC, which specialized in integrating heterogeneous information bases. However, recent emerging technologies such as internetworking and the World Wide Web have significantly expanded the types, availability, and volume of data available to a ...

6 **Dynamic service matchmaking among agents in open information environments**

Katia Sycara, Matthias Klusch, Seth Widoff, Jianguo Lu
March 1999 **ACM SIGMOD Record**, Volume 28 Issue 1

Full text available: [pdf\(702.54 KB\)](#) Additional Information: [full citation](#), [citations](#), [index terms](#)

7 **Putting it together: Internet privacy: a public concern**

Lorrie Faith Cranor
June 1998 **netWorker**, Volume 2 Issue 3

Full text available: [pdf\(336.26 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

8 **Privacy online**

Herman T. Tavani
December 1999 **ACM SIGCAS Computers and Society**, Volume 29 Issue 4

Full text available: [pdf\(1.06 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#)

9 **Augmenting paper to enhance community information sharing**

Antonietta Grasso, Alain Karsenty, Marco Susani
April 2000 **Proceedings of DARE 2000 on Designing augmented reality environments**

Full text available: [pdf\(903.95 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Paper is traditionally considered as a major gap between the physical and electronic worlds, especially after the many attempts that have failed to attain a completely digital world. Paper based artifacts have many affordances that people want to continue to exploit.

The work presented here is part of the Campiello project. It describes how the existing paper artifacts in use during the visits to cultural and tourist towns as well as artefacts used for local communities can be extended ...

Keywords: augmented reality, paper user interface, recommender systems

10 WebMate: a personal agent for browsing and searching

Liren Chen, Katia Sycara

May 1998 **Proceedings of the second international conference on Autonomous agents**

Full text available: [pdf\(969.48 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

11 Digital newspapers explore marketing on the Internet

Jonathan W. Palmer, Lars Bo Eriksen

September 1999 **Communications of the ACM**, Volume 42 Issue 9

Full text available: [pdf\(398.18 KB\)](#) [html\(34.48 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

12 Polychannel systems for mass digital communications

David K. Gifford

February 1990 **Communications of the ACM**, Volume 33 Issue 2

Full text available: [pdf\(2.22 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

A new type of distributed computer system looks toward the time when computers will be used by everyone everywhere.

13 Marketing information on the I-Way: data junkyard or information gold mine?

P. K. Kannan, Ai-Mei Chang, Andrew B. Whinston

March 1998 **Communications of the ACM**, Volume 41 Issue 3

Full text available: [pdf\(181.31 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

14 Social information filtering: algorithms for automating "word of mouth"

Upendra Shardanand, Pattie Maes

May 1995 **Proceedings of the SIGCHI conference on Human factors in computing systems**

Full text available: [html\(37.48 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

15 Building interfaces as personal agents: a case study

Amedeo Cesta, Daniel D'Aloisi

July 1996 **ACM SIGCHI Bulletin**, Volume 28 Issue 3

Full text available: [pdf\(725.61 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

This paper concerns the development of interfaces which perform tasks on behalf of the user. Recently the concept of task delegation has gained consideration due to the increasing number of assignments that are quite repetitive and tedious, like dealing with electronic messages, managing personal agendas, retrieving data and information in remote and distributed repositories. These tasks are generally connected to the development of computer networks. The design of such interfaces presents a hig ...

16 Survey articles: Web usage mining: discovery and applications of usage patterns from Web data

Jaideep Srivastava, Robert Cooley, Mukund Deshpande, Pang-Ning Tan
January 2000 **ACM SIGKDD Explorations Newsletter**, Volume 1 Issue 2

Full text available:  pdf(1.44 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Web usage mining is the application of data mining techniques to discover usage patterns from Web data, in order to understand and better serve the needs of Web-based applications. Web usage mining consists of three phases, namely *preprocessing*, *pattern discovery*, and *pattern analysis*. This paper describes each of these phases in detail. Given its application potential, Web usage mining has seen a rapid increase in interest, from both the research and practice communities. This pap ...

Keywords: data mining, web usage mining, world wide web

17 Varying the user interaction within multi-agent systems

Terry R. Payne, Katia Sycara, Michael Lewis
June 2000 **Proceedings of the fourth international conference on Autonomous agents**

Full text available:  pdf(934.38 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: collaboration, human-agent interaction, interface agents, matchmakers, middle agents, multi-agent teams, task agents

18 The personal electronic program guide—towards the pre-selection of individual TV programs

Michael Ehrmantraut, Theo Härdter, Hartmut Wittig, Ralf Steinmetz
November 1996 **Proceedings of the fifth international conference on Information and knowledge management**

Full text available:  pdf(923.87 KB) Additional Information: [full citation](#), [references](#), [index terms](#)

19 Conference review

Stuart Lowry
September 1999 **intelligence**, Volume 10 Issue 3

Full text available:  pdf(184.05 KB)  html(19.75 KB) Additional Information: [full citation](#), [index terms](#)

20 Staffing user services

A. Faye Borthick
November 1975 **Proceedings of the 3rd annual ACM SIGUCCS conference on User services**

Full text available:  pdf(242.01 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Staffing the User Services function of a computing center requires recruiting, training, and keeping qualified personnel while expending a fixed set of resources. These resources include not only the purely monetary ones of salaries and physical facilities but also other resources such as supervisory time. This paper addresses these activities from the viewpoint of how to plan the staffing of the User Services function given that one is prepared to delineate the objectives of the User Servi ...

Results 1 - 20 of 104

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

[IEEE HOME](#) | [SEARCH IEEE](#) | [SHOP](#) | [WEB ACCOUNT](#) | [CONTACT IEEE](#)[Membership](#) [Publications/Services](#) [Standards](#) [Conferences](#) [Careers/Jobs](#)

RELEASE 1.8

Welcome
United States Patent and Trademark Office

» Se.

[Help](#) [FAQ](#) [Terms](#) [IEEE Peer Review](#)**Quick Links** **Welcome to IEEE Xplore®**

- [Home](#)
- [What Can I Access?](#)
- [Log-out](#)

Tables of Contents

- [Journals & Magazines](#)
- [Conference Proceedings](#)
- [Standards](#)

Search

- [By Author](#)
- [Basic](#)
- [Advanced](#)
- [CrossRef](#)

Member Services

- [Join IEEE](#)
- [Establish IEEE Web Account](#)
- [Access the IEEE Member Digital Library](#)

IEEE Enterprise

- [Access the IEEE Enterprise File Cabinet](#)

 Print Format[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#) | [Join IEEE](#) | [Web Account](#) | [New this week](#) | [OPAC Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#) | [No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)

Copyright © 2004 IEEE — All rights reserved

[IEEE HOME](#) | [SEARCH IEEE](#) | [SHOP](#) | [WEB ACCOUNT](#) | [CONTACT IEEE](#)[Membership](#) [Publications/Services](#) [Standards](#) [Conferences](#) [Careers/Jobs](#)**IEEE Xplore®**
RELEASE 1.8Welcome
United States Patent and Trademark Office[Help](#) [FAQ](#) [Terms](#) [IEEE Peer Review](#)**Quick Links****Welcome to IEEE Xplore®**

- Home
- What Can I Access?
- Log-out

Tables of Contents

- Journals & Magazines
- Conference Proceedings
- Standards

Search

- By Author
- Basic
- Advanced
- CrossRef

Member Services

- Join IEEE
- Establish IEEE Web Account
- Access the IEEE Member Digital Library

IEEE Enterprise

- Access the IEEE Enterprise File Cabinet

Print Format

[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#) | [Join IEEE](#) | [Web Account](#) | [New this week](#) | [OPAC Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#) | [No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)

Copyright © 2004 IEEE — All rights reserved

An Intelligent Agent for Web Advertisements

Vincent Ng¹ and Kwan-Ho Mok

Department of Computing, Hong Kong Polytechnic University,
Kowloon, Hong Kong, China

Abstract

The rapid growth of Internet users attracts advertisers to post their advertisements in Internet. The probabilistic selection algorithm was not satisfactory; while other advertising agents are unable to guarantee the quality due to insufficient and unstable user information. This paper describes a new advertising agent based on user information. The users' interests are discovered by the Order Pattern Mining algorithm first, then applied the Gaussian curve transformation to represent their profiles. For the advertisements, we use the keywords from different categories to construct the advertisement profiles as Gaussian curves also. This allows us to select advertisements based on the intersections of the different profiles according to the users' preferences in an effective and efficient mechanism. A prototype of the Intelligent Advertising Agent has been developed with Java and Oracle. From our evaluations, we observed that about 80% of the test cases are successful in making predictions which generated the most favorable category that the users are interested.

1. Introduction

With the rapid growth of Internet users, the World Wide Web is a good way of presenting information to the public. The population connected to the Internet will grow from 30 million to more than 200 million by the year 2000 according to Input, a global IT market intelligence firm. More than 30,000 business information were posted in Internet and the number is expected to double in 1998 [1]. However, the advertising mode has been kept unchange which is similar to that used in TV and newspapers in essence.

On the Internet, the advertisements usually are in the form of clickable banners. When a user clicks on the banner, the web browser contacts the web server and returns the URL address of designated advertiser. The URL address could either be the home pages of the advertiser or a

¹The work of the authors were supported in part by the Hong Kong RGC Grant: PolyU 87/96E.

special site which encourages potential customers to visit and to get more detail information. Obviously, Internet advertising is more attractive than traditional methods because of the much lower cost. Further, advertising providers can have operations [2] such as

- determine the return on investments of their on-line advertising dollars
- how many customers have clicked on the advertisements
- on-line, centralised and real time reporting
- real time banner substitutions
- sophisticated tracking and reporting such as clicked rate and etc.
- comprehensive advertisement targeting capabilities on the Internet such as the browser type, search keywords and etc.

Although an advertising provider has the potential to give the above sophisticated features to advertisers, most current systems do not consider the interests of their users and will only display advertisements according to their own ideas of users' interests. In this paper, we suggest a new approach for creating an intelligent agent for advertising in Internet.

This paper is organised into six sections. Section 2 discusses some of the recent approaches in Internet advertisements and information filtering. Section 3 discusses our design of the Intelligent Advertising Agent (IAA). Section 4 will cover the implementation of the agent. Section 5 presents our experiments on the IAA prototype and Section 6 concludes our work and discusses the possible enhancements.

2. Recent Work

One basic paradigm mentioned in [4,5,6] for Internet advertising decisions is an uninformed approach which selects

an advertisement according to a probability based on the number of purchased units. Its problems has been identified as follows

- A user may see advertisements that are unrelated to his current interests.
- A user wastes a lot of time and money to download the uninterested information.
- Unexpected advertisements would irritate users in much the same way as a magazine article is split up with intervening advertisements.
- When a user accesses a Web page at different times, same advertisements are shown. It does not only waste the downloading time but also reduce the effectiveness of the advertisement.

One company, DoubleClick [7], is the first to offer advertisers the ability to dynamically target advertisements on the Web. They developed the DART targeting technology to offer four basic categories of targeting criteria: content targeting, behavioral targeting, user targeting, technical targeting. For DoubleClick or other similar system, the advertisement selection process is no different from a typical information filtering application. The idea is to select the information that a user prefers, or according to a user profile.

One approach in content-based filtering, a document representation can be derived from the document contents. Yan implemented a simple content-based text filtering system for Internet News articles in a system called 'SIFT' [9]. Profiles for SIFT are constructed manually by specifying words *prefer* or *avoid*. SIFT offers two facilities to assist users with profile constructions. Users are initially offered an opportunity to apply candidate profiles against the current day's articles to determine whether appropriate sets of articles are accepted and rejected. To facilitate maintenance of profiles over time, words which contributed to the position of each article in the ranked list are highlighted when using a web browser to access the articles. By examining the context of words with meanings that were unforeseen at the time the profile was constructed, users can select additional words which appear in the same context to add to the lists of words to be avoided.

Stevens, adopting the implicit acquisition, developed a system call InfoScope, also assigned to Internet news. It used automatic profile learning to minimise the complexity of exploiting information about the context in which words were used [11]. InfoScope deduced exact-match rules and offered them for approval by the users. These suggestions were based on simple observable actions such as the time spent reading a newsgroups or whether an individual message was saved for future reference to determine the user's interests. A user profile can be built up for information filtering.

There are limitations in InfoScope. The experimental system Stevens developed is able to process only information in the header of each article (e.g. subject, author, or newsgroup). Another explicit acquisition model has been developed by Alfred Kobsa, Andreas Nill, and Josef Fink jointly. They use the KN-AHS with BGP-MS [12,10] to demonstrate the feasibility of the user modeling for information filtering. The BGP-MS is an adaptive user modeling shell system. It utilises the partition mechanism SB-PART [11] which allows different types of assumptions about a user to represented simultaneously, but still separately. These assumptions include concerning a user's knowledge or goals, application-relevant characteristics of user subgroups (so-called 'stereotypes') or domain knowledge of the user modeling component.

With the exception of InfoScope, every system we have described requires a user to explicitly evaluate documents. Explicit feedback has the advantage of simplicity, and can minimise the source of experimental error which inference of the user's true reaction. Due to the insufficient capability of the existing algorithms, we decide to take this opportunity to propose a new approach to tackle the advertisement selection. The new solution will not depend on a user's explicit input. The agent will automatically extract the personal information, and based on it to select advertisements might be interested by them.

3. Designing the Internet Advertisement Agent

In designing the Intelligent Advertisement Agent (IAA), we have adopted the content-based filtering approach. We observe that there are several types of information available in Internet but the major information is still text. Selecting an advertisement from a list of advertisements can be considered as an information retrieval process. As Belkin and Croft observed, content-based text selection techniques have been extensively evaluated in the context of information retrieval [13]. In the following sections, we will discuss the design of IAA with respect to the four components in details.

3.1. Advertisement Context Representation

As discussed previously, we will use the content of advertisements to match against what information has been read by a visitor. In retrieving text information, it is usually in the form of keywords or in word phrases. The advantage in word phrase is that the meaning of the word phrase is very precise, thus reducing ambiguity in index and searching. However, the disadvantage is that we have to be aware of the phrase construction rules employed. The knowledge base can adopt an intermediate way by allowing both phrases and single words. This type of knowledge

Category	Number of keywords
AUTOMOBILE	43
COMPUTER	58
ENTERTAINMENT	49
FASHION	45
FINANCE	69
FOOD	31
HIFI	16
HOME AFFAIR	31
INTERNET	36
PHOTOGRAPHY	10
PROPERTY	62
SOFTWARE	26
SPORT	5
TELECOM	25
TRAINING	58
TRANSPORT	24
TRAVEL	118
WATCH	27

Figure 1. The number of keywords in each advertisement category.

base is usually manually constructed as automatic phrase construction is still difficult [8,15]. Furthermore, this can guarantee the specificity of keywords in representing their categories. In our design of IAA, we use the intermediate approach.

3.1.1 Advertisement Category

It is hard to decide the possible advertisement categories. If there are too few categories, they cannot clearly distinct users' interests; while too many categories will result with very few keywords in each of them. Initially, we adopt the advertisement categories from the classified section of a local newspaper, the *South China Morning Post*. At the same time, we also combine these advertisement categories with those found in Internet directories such as [3,19]. Figure 1 shows the different categories used in our final set.

3.1.2 Weighting Factor for Keywords

Another piece of information for content-based text filtering is the term-frequency which counts the number of times a keyword appears. The term-frequency can also be used to represent a keyword's importance. One technique is to represent each advertisement as a vector of TFIDF (Term Frequency Inverse Document Frequency) in the space of words that appeared in a set of training advertisements [14]. However, in a text filtering system, advance knowledge of the number of advertisements with a particular term is clearly not possible. Estimates are based on sampling earlier advertisements to produce useful inverse document frequency values for domains in which term usage patterns are relatively stable [8]. In our design, we adopt a simple algorithm to generate the weighting factor. The weighting factor for a keyword in each advertisement category is calculated by the following equation:

$$WeightFactor_i = \frac{F_i}{T} \quad (1)$$

where F_i is number of occurrences of term in N_k advertisements of category k , and T is the total occurrences of all terms in N_k advertisements in the same category.

3.2. Profile Construction

In an information filtering system, the system's representation of the information needed is commonly referred to as a 'profile'. It would not be technically correct to call the profile a 'user model' because the user model consists of both a representation of a user and a method for interpreting that representation to make predictions. Here, we shall use a content-based filtering with a new model for user and advertisement profiles to achieve this task.

3.2.1 Gaussian Model

Many personalized systems uses a single scalar to rank a document's interest for a user such as in [16]. Usually, a high value indicates a high level of interest whereas a low value is unimportant to a user. However, this approach does not have the broadness of a user's interest. A better technique, as proposed in [17], is to use two values. One value is to focus on the median (μ) of the user's interest in a scale from *highly interested* to *totally dislikes*, and a second value to describe the broadness (δ) of the user's interest [17].

3.2.2 Gaussian Curve Model for Advertisement and User

We can imagine that a user may have different interests on different advertisement categories. Therefore a user profile should include several Gaussian curves to represent a user's interests on different categories. Each Gaussian curve for a category in a user profile is called *category-profile*. That is, if there are m advertisement categories, there are m category-profile for the user. With a Gaussian curve, we can roughly represent three levels of interests:

- Dislike the advertisement
- Indifferent to the advertisement
- High interest the advertisement

It is possible to have nine combinations by mapping the user and advertisement profiles. Yet, when an advertisement belongs to a certain category, it should have a high score already. Therefore, the possible mapping can be reduced to three.

3.2.3 Advertisement Category Profiles

When we read advertisements, we always find the same keyword frequently appears for different advertisements in

the same category. Furthermore, some advertisements may fit in different categories as well. This means that one keyword can appear in different categories and one advertisement can belong to more than one category. Therefore, rather than use distinct sets of keywords for different advertisement categories, keywords can be used for several categories except they may have different weights.

An advertisement curve for each category represents the characteristics of the advertisements in that category. When advertisers want to show their advertisements in the Web pages of an ISP (Internet Service Provider), they will make contact with the ISP. An advertiser will specify his target groups and allocate his advertisements in the matching categories. To identify an advertisement's characteristics, keywords in the advertisement will be extracted. The extracted keywords will be matched with those in the same category in the knowledge base. Each keyword in the knowledge base is associated with a weighting factor which is discussed in Section 3.1.2. When keywords are matched, the distribution of their weights will represent the characteristic of the advertisement. After processing of all advertisements in the same category, we use a Gaussian curve (normal distribution curve) to represent its profile [18].

3.2.4 User Profiles

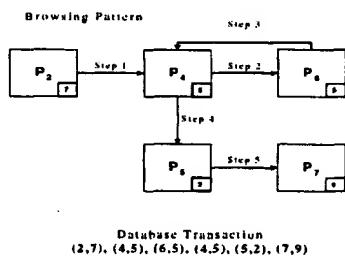


Figure 2. An access pattern.

The concept of user profiles is similar to that of an advertisement. The difference is that the keyword extractions is from the visited URLs of a user. It is because the content of the URLs (uniform resource locators) is indirectly representing a user's interests. For most WWW servers, there are access logs to keep track of their utilizations. Each line in the log file indicates one access of a particular Web page. This is only interesting from a single server view, but not for individual users. In fact, we are interested in the access patterns of web pages that either their total viewing times are

more than a threshold or the patterns are having a certain or more number of Web pages. Here, we adopt the OPM (Order Pattern Mining) algorithm in discovering the access patterns that will indicate users' interests [20].

In order to work with the OPM algorithm, we need to identify Web pages that are accessed by an individual user, or even a session done by an individual user. A normal WWW access log records the information of client IP address, name of html pages and the access time on those pages. With these information, our identification method is based on the followings assumptions

- The server can be a WWW server or a proxy server as long as access logs can be recorded.
- Only one user is active on a single IP address at a time.
- When a user has been idle for a long time (say 1 hour), the current session ends.
- Local cache in the client machines have been disabled and every Web page access will need to be served by the server.

Each session of a user is transformed into an *access transaction*. For example, in Figure 2, User A accessed the Web pages in the sequence of $P_2 \rightarrow P_4 \rightarrow P_6 \rightarrow P_4 \rightarrow P_5 \rightarrow P_7$. The access sequence forms a list of order pairs in the format of (P_i, d_{ij}) where P_i is the page label (number) and d_{ij} the viewing time of the page. Each order pair of T_i is represented as W_{ij} . For the example in Figure 2, the sequence of the access transaction is $\{(2,7), (4,5), (6,5), (4,5), (5,2), (7,9)\}$.

From the access log, we identify the sessions, convert them into access transactions and then prepare for mining of user behaviors. After these transformation, we can apply the OPM algorithm to solve the following problem.

Problem 1 For all access transactions in a given database D , extract all the frequent patterns of size k (AP_k 's) for different users where for each equivalent access pattern S_k^i of a AP_k , we have

$$\sum d_{ij} \geq F.$$

F is a threshold value representing the minimal viewing time of the pattern.

The extracted access patterns reflect the interest of users. Keywords from the web pages of the access patterns of a user will be matched with keywords with each advertisement category. After the completion of the matching process with all advertisement categories, we construct the user category profiles. For example, a user profile of a category is represented by the Gaussian curve in Figure 3 with $\mu = 0.3$ and $\delta = 0.5$.

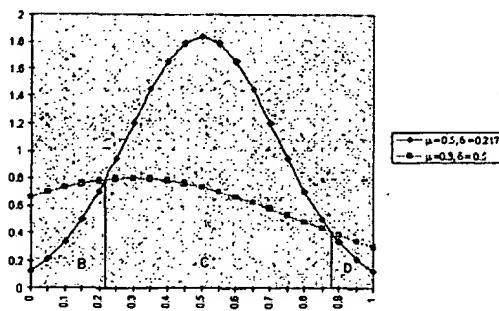


Figure 3. Calculation of overlapping area by portion.

3.3. Profile Comparisons

From last two sections, Gaussian curves are used to model a user's interests and an advertisement's characteristics. Therefore, we can decide a selection by measuring the amount of overlaps. That is, if the overlapping area is large, the correlation between two profiles is relatively close and vice versa. An example for the comparison is in Figure 3.

Obviously, the areas under two Gaussian curves may have different sizes and shapes. One calculation method is to discretize the parameters and store the results in a table for quick retrievals later. The first step is to transform all the observations of a normal random variable X into a normal variable Z with mean zero and variance one. This can be done by the transformation

$$Z = (X - \mu)/\delta \quad (2)$$

To calculate the overlapping area of two curves, we find the intersection points with the given μ_1, δ_1, μ_2 and δ_2 by using the formula below.

$$\frac{1}{\sqrt{2\pi}\delta_1} e^{\frac{1}{2}(\frac{x-\mu_1}{\delta_1})^2} = \frac{1}{\sqrt{2\pi}\delta_2} e^{\frac{1}{2}(\frac{x-\mu_2}{\delta_2})^2} \quad (3)$$

For the example in Figure 4, the X values of the intersection points are 0.217 and 0.875. Their corresponding values of Z_1 and Z_2 of the curve with square symbols are equal to

$$Z_1 = (0.217 - 0.3)/0.5 = -0.166 \quad (4)$$

and

$$Z_2 = (0.875 - 0.3)/0.5 = 1.15 \quad (5)$$

The corresponding values of Z_3 and Z_4 of the second curve (with diamond symbols) are equal to -1.304 and 1.728 respectively. From the Z-value table in [18], the corresponding areas (A_1, A_2, A_3, A_4) for Z_1, Z_2, Z_3 and Z_4 are

1. Let U_{p1}, \dots, U_{pk} be the user's category profiles (Gaussian curves) of user U .
2. For all advertisement categories (C_1, \dots, C_k), find the one category (C_i) that has the maximum overlapping area from $C_i \cap U_{pi}$.
3. For all advertisements in C_i , let A_{i1}, \dots, A_{i5} be the five advertisements with the highest scores calculated by
 - score = OverlapArea(U_p, A_{ij}) \times NU_{ij}
where NU_{ij} is the percentage of unused units of advertisement A_{ij} over the sum of all unused units in advertisement category i .
4. Randomly select one of the five advertisements for the user. Decrease the number of unused of that advertisement by one.

Figure 4. Algorithm to select an advertisement.

0.4325, 0.8749, 0.0968 and 0.9528, respectively. Therefore, the overlapping area of the two curves can be found as the sum of areas 'B', 'C' and 'D' in the figure, and it is equal to

$$Total Area = A_1 + (A_4 - A_3) + (1 - A_2) = 0.5864. \quad (6)$$

3.4. Selecting an Advertisement

We can now derive a measure to indicate if a user will accept an advertisement category for his viewing. In the IAA, an advertisement selection algorithm is developed as shown in Figure 4. It considers the advertisement characteristics and the display units bought for an advertisement. In the algorithm, we use the product of the advertisement score (overlapping area) with the percentage of the unused display units for all advertisements in the selected category. From the five highest scored advertisements, we randomly select one for the user. In this way, it tries to prevent from 'banner burnout' and maintains healthy click-through rates.

4. IAA Development

The system architecture for the intelligent advertising agent (IAA) is proposed in Figure 5. It has four main modules:

- Text Retrieval Module
- Profile Modeling Module
- Advertisement User Matching Module

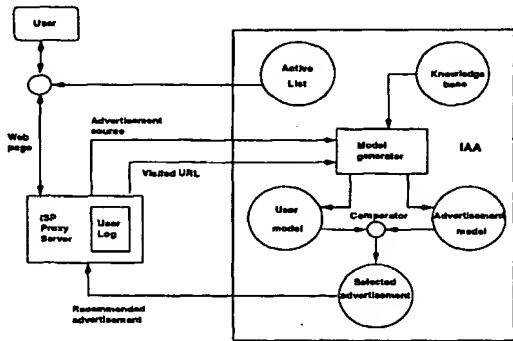


Figure 5. IAA Architecture.

- Advertisement Display Module

A prototype of the IAA is developed with Oracle and Java. During the implementation, we want to reduce our development effort while the prototype should have sufficient features to demonstrate the effectiveness and accuracy of our IAA design. Therefore, we have relaxed our requirements with the followings.

- The prototype is to be tested offline. That is, advertisements will not be shown during regular web browsing.
- User web browsing histories (preferences) are collected either via a proxy server or manually. This is to support our user profile creations.
- Advertisements and their contents are acquired either via the popular Internet gateways or manually.

The knowledge base is the key area in the system. The keywords in it are used to generate the user and advertisement profiles. In this pre-processing steps, a set of advertisements (URLs) are collected and their profiles are generated for the later matching process. In calculating the overlappings of the profiles, we use the one-byte representation scheme as in [17]. A single byte is used to pack the two parameters, μ and δ , of a profile after discretizing into a range from 0 to 255. Pre-calculated values are stored in the Gaussian lookup table and the Z-value table. This would allow us to find out the overlapping area by a simple 3-way join over the tables.

5. Evaluation

In our experiments, we would like to measure the predictive performance of the IAA prototype by the acceptance of the suggested advertisements for our users. In doing the experiments, we focused on three main aspects: keyword extraction, matching model for predictions and the accuracy of the pre-calculated tables.

Source	Case Number	No. of users
PolyU students	1	1
Colleagues	2-7	6
PolyU students	8-14	27
Secondary students	35-85	54
Other universities students	89-128	40

Figure 6. Users evaluating IAA.

Advert. Cat.	No. of Advert.	μ^*	σ^*	Packed μ^* and σ^*
Computer	7	0.1333	0.0667	.94
Software	13	0.2	0.1333	.77
Internet	8	0.467	0.0667	.29
Hi-Fi	3	0.1333	0.0667	.94
Finance	5	0.4	0.2556	.27

Figure 7. Advertisements to evaluate IAA.

At the beginning, after we have completed our first IAA prototype, we invited a group of students in Hong Kong Polytechnic University and colleagues of the two authors. Unfortunately, there were only 7 responses. The poor response may be due to the timing, when the invitations were done close to the end of the term. We then tried again in next study term but also with other sources, secondary schools and other universities. The final distribution of users evaluating the IAA prototype is shown in Figure 6. We believe that our test cases are more computer literate. Also, because of limited resources, we decided to use only 5 advertisement categories, which have some relationship to computing, in our experiments. Advertisements for these five categories have been extracted via the Internet directory home pages [3,19]. Their categories are set to the same categories as originally placed in the Internet directory. The advertisement categories and their profiles are shown in Figure 7.

5.1. User Profiles

There are two batches of our cases. The first batch is of only 7 users. For this batch of users, we asked them to fill in a questionnaire before trying IAA. The questionnaire has two sections. The first section will ask a user to provide 18 web pages that he usually visits. The second section is to assign weights to the 5 advertisement categories to verify the advertisements that the users are usually interested. After receiving the survey forms, we divided the advertisements into two groups. One was the targeted advertisement categories including computer hardware, software and Internet. The second one was the uninterested advertisement categories including finance and audio.

For the second batch, we obtain the access logs from their web masters with the permissions of the users. The OPM algorithm is applied to discover the access patterns and hence the web pages for those users visit frequently. As in the first batch, we also ask the users to rank the advertisement category manually.

For all users in our experiments, user profiles are con-

Case Name	Category	μ'	σ'	Packed μ' and σ'
Case 1	Computer	0.1333	0.0667	-94
Case 1	Internet	0.3335	0.267	-43
Case 1	Software	0.2	0.1333	-77
Case 1	Hi-Fi	0	0	-127
Case 1	Finance	0.0667	0.0667	-110
Case 2	Computer	0.2	0.1333	-93
Case 2	Internet	0.0667	0.0667	-110
Case 2	Software	0.267	0.1333	-77
Case 2	Hi-Fi	0.0667	0.0667	-110
Case 2	Finance	0	0	-127
Case 3	Computer	0.2	0.1333	-93
Case 3	Internet	0.2667	0.2667	-59
Case 3	Software	0.267	0.1333	-77
Case 3	Hi-Fi	0	0	-127
Case 3	Finance	0.0667	0.0667	-110

Figure 8. Generated user profiles.

structured according to the content of the web pages found. As there are five advertisement categories in our experiments, only the profiles belonged to the five categories are listed out. The μ 's and σ 's of the user category profiles for the first 3 users are shown in Figure 8.

5.2. Prediction Accuracy

Each user profile is matched with the advertisement profiles to determine what advertisement categories to be selected. The matching results of 4 different groups of the test cases are shown in Figure 9. The advertisement categories recommended by the IAA to users are marked with '*'*. '1' represents the totally overlapped, whereas '0' represents close to zero.

In order to improve the performance of calculating the overlapping areas, all possible discrete values of two Gaussian curves are saved in a database table. In our experiments, we compared the results of the approximations with the exact calculations as shown in Figure 10. We observed that there are discrepancies, such as for Case 7. We believe this is due to round-off errors as single bytes are used in the pre-calculated tables. Instead of asking the user to verify the selections, we compare the order of the five selected categories (the 5 largest overlapping areas) with the same five categories ranked by the user in the questionnaire. Suppose the five selected categories are C_1, \dots, C_5 and the ranked categories are R_1, \dots, R_5 . We measure the prediction correctness as

$$P_c = \frac{1}{5} \sum_{i=1}^5 (C_i = R_i) \quad (7)$$

When P_c is 1, the prediction is perfect. When it is 0, the prediction is completely inaccurate.

Case Name	Category	Approximated Result	Exact Result
Case 1	Computer	0.802	0.756
Case 1	Internet	0.6603	0.6766
Case 1	Software	1*	0.9034*
Case 1	Hi-Fi	0	0.3815
Case 1	Finance	0.2236	0.2388
Case 2	Computer	0.6772	0.9407*
Case 2	Internet	0.0898	0.0418
Case 2	Software	1*	0.9034
Case 2	Hi-Fi	0.617	0.325
Case 2	Finance	0	0.0442
Case 3	Computer	0.6772	0.9189
Case 3	Internet	0.6093	0.6842
Case 3	Software	1*	0.9213*
Case 3	Hi-Fi	0	0.3827
Case 3	Finance	0.2236	0.3007
Case 7	Computer	1*	0.7766
Case 7	Internet	0.7623	0.7573
Case 7	Software	0.8026	0.9058*
Case 7	Hi-Fi	0	0.2313
Case 7	Finance	0.4536	0.4557

Figure 9. Partial results of profile overlapping.

5.3. IAA Prototype Performance

From our evaluations, we observed that about 16% of our cases have discrepancies between exact and approximated calculations. Therefore, we found that the precision of the pre-calculation overlapping area technique is insufficient. If the precision is improved, the prototype can produce more accurate predictions. To do that, the discrete values of the Gaussian curve should use double bytes instead of single byte. Figure 10 shows the prediction performance of the IAA prototype using the exact results. Overall, the predictions are doing quite well except for the secondary students. We believe that secondary students have a wider scope interest and are not as focused as computing professionals or to-be professionals.

Here, we would like to discuss some sample cases which can be generalized for others. For example, we found that the results of Case 1 and Case 3 are very close to the user's expectation. The recommended category for Case 1 is software which matches with the most interested category of the user. In addition, the results of computer, software and Internet categories have large overlapping areas. At the same time, IAA can also filter out the uninterested categories such as the finance and Hi-Fi.

Another example is Case 2 where we found that IAA recommends the Software category with approximated result but the correct category for Case 2 with the exact result. In the case, we found that Hi-Fi category has a relative large overlapping area. It is because the number of key-

Source	Average P_c
PolyU students	0.74
Colleagues	0.86
Senior Secondary students	0.62
Other universities	0.84

Figure 10. Average P_c of users evaluating the IAA prototype.

words in Hi-Fi was too small and the advertisement profile was close to zero. Under this circumstance, a misled result is produced. In our design, we used the advertisement profile as a control to prevent from falling in the *dislike* situation. Therefore, the overlapping area of Hi-Fi can be reduced if the profile has stronger category characteristics, far from zero.

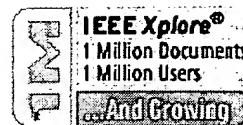
6. Conclusion

The rapid growth of Internet users attracts advertisers to post their advertisements in Internet. The probabilistic selection algorithm was not satisfactory [4]; while other advertising agents are unable to guarantee the quality due to insufficient and unstable user information. We took this opportunity to develop a new advertising agent based on user information. By using the keyword knowledge base and Gaussian curve transformation, we can determine users' interests when comparing the advertisement and user profiles in an effective and efficient mechanism.

A prototype of the Intelligent Advertising Agent has been developed with Java and Oracle. From our evaluations, we observed that about 80% of the test cases are successful in making predictions which generated the most favorable category that the users are interested.

References

- [1] <http://www.iwcom.com/presence.htm>
- [2] <http://www.gobeyond.com>.
- [3] Yahoo!, Internet Directory, <http://www.yahoo.com>.
- [4] Yager, R. Ronald, 1997, Intelligent Agents for World Wide Web Advertising Decisions. International Journal of Intelligent Systems. Vol. 12, pp 379-390.
- [5] Zeff R. and Aronson B., 1997, Learning About Your Customer and Delivery Targeted Ads, Web Advertising '97, pp 1-4.
- [6] http://www.cookiecentral.com/c_concept.htm.
- [7] Double Click, Internet Advertising Solution Company, <http://www.doubleclick.net>.
- [8] Oard W. D. and Marchionini G. 1996, A Conceptual Framework for Text Filtering. University of Maryland, Technical Report EE-TR-96-25 CAR-TR-830 CLIS-TR-96-02 CS-TR-3643, pp 1-32.
- [9] Yan W.T. and Molina G.H. 1994, Distributed selective dissemination of Information. In proceedings of the Third International Conference on Parallel and Distributed Information Systems, pp 89-98. IEEE Computer Society.
- [10] Kobsa, A. 1990, Modelling the User's Conceptual Knowledge in BGP-MS, a User Modelling shell System. Computational Intelligence Vol. 6, pp 193-208.
- [11] Kobsa, A., Nill A., and Fink J. 1994, Adaptive Hypertext and Hypermedia Clients of the User Modeling System BGP-MS, pp 1-23.
- [12] Rich. E.A. 1979, User Modelling via Stereotypes. Cognitive Science, Vol. 3, pp 329-354.
- [13] Belkin J. N. and Croft B.W. 1992, Information filtering and Information Retrieval: Two sides of the same coin? Communications of the ACM, Vol. 35 No. 12, pp 29-38.
- [14] Salton G. 1987, Term Weighting approaches in Automatic Text Retrieval, Technical report, COR-87-881, Department of Computer Science, Cornell University, pp 1-23.
- [15] Frakes B. W. and Yates B. R. 1992, Information Retrieval : Data Structures & Algorithms, Prentice Hall PTR, pp 102-160.
- [16] Goldberg D. et al. 1992, Using Collaborative Filtering to Weave an Information Tapestry, Communications of the ACM, 1992, Vol. 25 No. 12, pp 61-70.
- [17] Kahabka T. and Mari K. and Specht G. 1997, GRAS: An Adaptive Personalization Scheme for Hypermedia Databases, Proc. of the Second Conference on Hypertext – Information Retrieval – Multimedia, pp 279-292.
- [18] Walpole R. 1982, Introduction to Statistics, Macmillan Publishing Co., Inc., pp 183-241.
- [19] Lycos!, Internet Directory, <http://www.lycos.com>.
- [20] Lo C and Ng V, 1999, Discovering Web Access Orders with Association Rules, IEEE SMC 1999, Vol 4, pp. 99-104.



Welcome to IEEE Xplore®

- Home
- What Can I Access?
- Log-out

Tables of Contents

- Journals & Magazines
- Conference Proceedings
- Standards

Search

- By Author
- Basic
- Advanced
- CrossRef

Member Services

- Join IEEE
- Establish IEEE Web Account
- Access the IEEE Member Digital Library

IEEE Enterprise

- Access the IEEE Enterprise File Cabinet

Print Format

Your search matched **40** of **1099723** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance Descending** order.

Refine This Search:

You may refine your search by editing the current search expression or enter a new one in the text box.

Check to search within this result set

Results Key:

JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

= Your access to full-text

1 An intelligent agent for Web advertisements

Ng, V.; Kwan-Ho Mok;

Cooperative Database Systems for Advanced Applications, 2001. CODAS 2001 Proceedings of the Third International Symposium on , 23-24 April 2001
Pages:102 - 109

[Abstract] [PDF Full-Text (596 KB)] IEEE CNF

2 Dynamic scheduling techniques for interactive hypermedia servers

Hamidzadeh, B.; Tsun-Ping, J.;

Consumer Electronics, IEEE Transactions on , Volume: 45 , Issue: 1 , Feb. 1999
Pages:46 - 56

[Abstract] [PDF Full-Text (908 KB)] IEEE JNL

3 Analysis of visual information recognition in Web site viewing

Ishikawa, T.; Kobayashi, K.; Watanabe, K.;

SICE 2003 Annual Conference , Volume: 1 , 4-6 Aug. 2003

Pages:58 - 61 Vol.1

[Abstract] [PDF Full-Text (338 KB)] IEEE CNF

4 The identification of missing information resources through the query difference operator

Minock, M.; Rusinkiewicz, M.; Perry, B.;

Cooperative Information Systems, 1999. CoopIS '99. Proceedings. 1999 IFCIS International Conference on , 2-4 Sept. 1999

Pages:304 - 313

[Abstract] [\[PDF Full-Text \(140 KB\)\]](#) IEEE CNF

5 Intelligent information recommendation system in the Internet

Jin-Gang Kim; Eun-Seok Lee;

Parallel Processing, 1999. Proceedings. 1999 International Workshops on , 21 Sept. 1999

Pages:376 - 380

[Abstract] [\[PDF Full-Text \(84 KB\)\]](#) IEEE CNF

6 SpaceTag: an overlaid virtual system and its applications

Tarumi, H.; Morishita, K.; Nakao, M.; Kambayashi, Y.;

Multimedia Computing and Systems, 1999. IEEE International Conference on , Volume: 1 , 7-11 June 1999

Pages:207 - 212 vol.1

[Abstract] [\[PDF Full-Text \(544 KB\)\]](#) IEEE CNF

7 Wireless advertising's challenges and opportunities

Mohamed Yunos, H.; Zeyu Gao, J.; Shim, S.;

Computer , Volume: 36 , Issue: 5 , May 2003

Pages:30 - 37

[Abstract] [\[PDF Full-Text \(2688 KB\)\]](#) IEEE JNL

8 QUARTS-II: a routing simulator for ATM networks

Sivabalan, M.; Mouftah, H.T.;

Communications Magazine, IEEE , Volume: 36 , Issue: 5 , May 1998

Pages:80 - 87

[Abstract] [\[PDF Full-Text \(1240 KB\)\]](#) IEEE JNL

9 Mobile IP

Perkins, C.E.;

Communications Magazine, IEEE , Volume: 35 , Issue: 5 , May 1997

Pages:84 - 99

[Abstract] [\[PDF Full-Text \(2472 KB\)\]](#) IEEE JNL

10 Hotring advertising model for broadband hotspots

Jamaluddin, J.; Doherty, M.; Edwards, R.; Coulton, P.;

Wireless Communications and Networking Conference, 2004. WCNC. 2004

IEEE , Volume: 4 , 21-25 March 2004

Pages:2064 - 2068 Vol.4

[Abstract] [\[PDF Full-Text \(315 KB\)\]](#) IEEE CNF

11 WHEELS: a conversational system in the automobile classifieds domain

Meng, H.; Busayapongchai, S.; Glass, J.; Goddeau, D.; Hethetingron, L.; Hurl E.; Pao, C.; Polifroni, J.; Seneff, S.; Zue, V.;

Spoken Language, 1996. ICSLP 96. Proceedings., Fourth International Conference on , Volume: 1 , 3-6 Oct. 1996

Pages:542 - 545 vol.1

[\[Abstract\]](#) [\[PDF Full-Text \(384 KB\)\]](#) [IEEE CNF](#)

12 A framework-based dialogue manager for spoken language applications
Goddeau, D.; Meng, H.; Polifroni, J.; Seneff, S.; Busayapongchai, S.;
Spoken Language, 1996. ICSLP 96. Proceedings., Fourth International Conference on , Volume: 2 , 3-6 Oct. 1996
Pages:701 - 704 vol.2

[\[Abstract\]](#) [\[PDF Full-Text \(392 KB\)\]](#) [IEEE CNF](#)

13 A system based on hierarchical fuzzy clustering for web users profile
Lazzerini, B.; Marcelloni, F.; Cococcioni, M.;
Systems, Man and Cybernetics, 2003. IEEE International Conference on , Volume 2 , 5-8 Oct. 2003
Pages:1995 - 2000 vol.2

[\[Abstract\]](#) [\[PDF Full-Text \(480 KB\)\]](#) [IEEE CNF](#)

14 SNAP&TELL: a multi-modal wearable computer interface for browsing the environment
Keaton, T.; Dominguez, S.M.; Sayed, A.H.;
Wearable Computers, 2002. (ISWC 2002). Proceedings. Sixth International Symposium on , 7-10 Oct. 2002
Pages:75 - 82

[\[Abstract\]](#) [\[PDF Full-Text \(2981 KB\)\]](#) [IEEE CNF](#)

15 XML-based household profiling for home networks
Kerlefsen, J.C.;
Consumer Electronics, 2001. ICCE. International Conference on , 19-21 June 2001
Pages:198 - 199

[\[Abstract\]](#) [\[PDF Full-Text \(144 KB\)\]](#) [IEEE CNF](#)

[1](#) [2](#) [3](#) [Next](#)

A Framework for Targeting Banner Advertising On the Internet

Katherine Gallagher and Jeffrey Parsons

Faculty of Business Administration, Memorial University of Newfoundland

St. John's, NF, Canada A1B 3X5

{kgallagh, jeffreyp}@morgan.ucs.mun.ca

Abstract

Constraints that limit accurate targeting of advertising in traditional media may not hold in cyberspace. This paper presents a model for effectively and efficiently targeting hypermedia-based banner advertisements in an online information service. The model takes advantage of information technology to micro-target banner advertisements based on individual characteristics of users. A simple version of the model, which has the virtue of ease of development, is presented. Enhancements are also proposed. These require more effort to develop, but may lead to even more precise targeting of advertisements. Implementation of this framework may benefit both online advertisers and online consumers.

1. Introduction

Cyberspace is a rapidly growing new medium for commerce. To date, a great deal of industry attention has focused on electronic transactions over the Internet. Although rapid growth is predicted over the next few years [10, 17, 21], actual sales thus far have been only moderate: users appear to regard the Internet primarily as a source of product information--when it comes time to pay, they prefer to buy offline by more conventional means [12, 14].

Responding to consumers' desire for information, businesses in large numbers have developed sites on the World Wide Web (WWW or Web). Most commercial Web sites describe the firm and its products and/or services, and many offer opportunities for visitors to the Web site to provide feedback and ask for specific information. As well, some Web sites collect information from visitors in order to improve future offerings. Some sites also support ordering and payment. The interactive potential of Web sites is particularly exciting, as it facilitates relationship marketing and customer support, eliminating the obstacles of geography and time [14, 22]. Not surprisingly, then, industry and scholarly research has recently focused on making Web sites more appealing and useful to visitors [13]. However, a Web site can only be effective if current and prospective customers visit it. Attracting this audience is currently a major challenge.

In this paper, we address the challenge of attracting a defined target audience to a Web site via *banner advertising*. We propose a framework for effectively targeting banner advertising in an electronic marketplace in a manner that benefits both advertisers and consumers. It allows advertisers to reach consumers who are more likely to be interested in the products and/or services offered by the company, and exposes consumers to information about products and services that they are likely to be interested in purchasing. Although the framework is discussed in terms of the Internet, we believe it will be relevant to whatever form the "information superhighway" eventually assumes. The framework takes advantage of the capabilities afforded by information technology for collecting and processing data about users. The next section examines trends in the electronic marketplace. Subsequently, the current state of advertising in this medium is discussed. Thereafter, a framework for targeting banner advertising, supported by appropriate information technologies, is proposed. Finally, opportunities for further research are discussed.

2. Marketing and Advertising in an Evolving Electronic Marketplace

The Internet began in the early 1970s as a US government research project designed primarily for the needs of the military. It expanded in the 1980s to serve the international academic and research communities [19, 23]. In the 1990s, businesses began to appear on the Internet. Although accurate estimates are obsolete as soon as they are made, it is clear that today tens of millions of people have access to the Internet [16] through over 100,000 computer networks in 150 countries--and the numbers continue to increase [14]. Two types of developments are particularly noteworthy with regard to this growth.

First, a large and ever expanding number of affluent, educated consumers are using the Internet [11]. This concentration of very desirable consumers has led to a surge in commercial interest. Prior to 1990, nodes on the Internet were predominantly academic institutions. In 1990, about 1,000 businesses had Internet connections. By June 1995, over 21,000 businesses were online, and the growth in commercial connectivity shows no sign of slowing [8].

Second, the emergence of the hypermedia-based WWW,

together with point-and-click multimedia interfaces such as Netscape, have greatly increased usability of the Internet for persons without extensive computer training. The development of "applet" technology, such as Java, which allows programs to run on a variety of platforms, increases the transparency of various Internet services. In other words, as technology continues to evolve, it is no longer an obstacle to, but an enabler of, electronic commerce.

In this environment, companies are seeking ways to use the Internet effectively [1, 3, 13, 22]. One active area in electronic commerce involves using the Internet as a medium to communicate persuasive product and service information via advertisements. These take various forms, the most common of which are corporate Web sites and banner advertising. We define a banner advertisement as:

- paid communication (via text, graphics, video and/or audio) of information about an organization and/or its products and services
- by an identified sponsor
- embedded within, and visually distinct from, information provided by an online service
- with hypermedia links to the sponsor's Web site.

We distinguish banner advertising from simple hypermedia links (paid or not) to commercial Web sites: banner advertising conveys a message even if the user does not follow the link; simple links can only convey a message if the user follows the link. Banner advertisements are also distinct from what [14] refer to as "flat ads," single page advertisements that do not contain hypermedia links. In this paper, we restrict our discussion to banner advertising that appears in the course of users' browsing and searching activities on information services, such as Yahoo! (<http://www.yahoo.com>) and Excite (<http://www.excite.com>), that provide an entry point to Internet resources. Appendix 1 shows a banner advertisements by the Saturn automobile company.

Scant attention has been paid to banner advertising by researchers. This may be because banners seem relatively insignificant, especially when compared with the interactive richness of Web sites. Technical specifications for banner advertisements severely limit creative options and preclude any consumer-firm interaction beyond the consumer's selection of the hypermedia link to the associated Web site (Excite, for instance, specifies that "all banners are 468x60 pixels, gif format only, maximum file size is 7k" [9]). Banner advertisements are, however, very important and interesting when viewed as part of a system that converts browsers and searchers into Web site visitors and, ultimately, customers. In their model of this conversion process, Berthon, Pitt and Watson [3] identify a sequence of tasks. First, users must be made aware of the Web site, then they must be attracted to and locate the site. Once users have found the Web site, the task is to turn that hit into a

visit, ensuring there is some meaningful contact between the firm and the consumer, then to convert the visit into a purchase. The final task is to get purchasers to return to the Web site and repurchase. Each task in the sequence is dependent on the successful execution of the previous task.

Our view of the role of banner advertising in this system is as a mechanism to make target audience members aware of a firm's Web site and to attract those users to the site. We define two concepts critical to understanding this role. *Attraction effectiveness* is the number of target audience members who reach a company's Web site via a banner advertisement hypermedia link divided by the number of target audience members who use the information service on which the advertisement appears. *Attraction efficiency* is the advertising cost per target audience member attracted to a company's Web site via a banner advertisement.

There is some evidence that the attraction efficiency of banner advertising is low. A recent estimate indicates that only 1-2% of banner advertisements lead viewers to seek additional information (e.g., by selecting a hypermedia link to the company's Web site) [5]. Since information services charge advertisers based on number of exposures (e.g., [9, 24]), the cost of attracting a single target audience member to a Web site is at least 50 to 100 times what it would be if all users who were exposed to the advertisement selected the hypermedia link. (The cost is even higher if some users selecting the link are not target audience members.) Increasing attraction efficiency by reducing wasted exposures should therefore be a priority. (An additional motivation for improving performance of banner advertising in converting searchers and browsers into Web site visitors arises from recent events such as the agreement between Procter & Gamble and Yahoo! which provides for payment based on the number of people who actually seek additional information (by selecting a link from a banner advertisement) rather than those who are merely exposed to the advertisement [20]. Such arrangements are expected to pressure online services to eliminate wasted exposures [5].)

The estimate cited above does not provide evidence on the attraction effectiveness of banner advertising. The fact that only 1-2% of exposed users select a link to the advertiser's site is irrelevant to effectiveness if all target audience members using the information service are among this group. However, since banner advertisements on online information services are shown selectively to users, there will generally be the possibility that some target audience members who use the information service will not be exposed to the advertisement and, hence, will be unable to link to a company's Web site via it. Depending on the strategy used to select advertisements for users, a large number of target audience members may be missed.

We contend that both the attraction effectiveness and efficiency of banner advertising can be improved by

precisely targeting advertisements based on characteristics and behavior of individual users of information services. Moreover, such targeting can be more precise than the targeting possible in traditional media. For example, visitors to a "Travel" page on an information service may be good targets for an advertisement for discount airfares, as would readers of the Travel section of a newspaper. But the fact that the online visitors have made a series of decisions and taken a series of actions (i.e., selecting only a subset of highlighted links within a hierarchical menu of categories) to reach the Travel page, rather than some other page (e.g., the Home Decorating page) suggests they may have a greater interest in travel than, say, readers who unintentionally come upon the Travel section of a newspaper and decide to read it. Since these exposures are more likely to be target audience members, attraction effectiveness can be improved. Targeting individual users strategy should also lead to fewer wasted exposures, since the advertisement would not be shown to users who have not reached the Travel page, thereby improving attraction effectiveness. (See Appendix 2 for a similar example.)

At present, targeting of banner advertising does not always occur. For example, Appendix 3 shows an advertisement for Honda that appeared when Organic Gardening was selected from a hierarchical menu of categories. People interested in organic gardening may not be the best prospects for automobiles, as they are likely to be more environmentally sensitive than the general population and may feel that cars unnecessarily harm the environment.

Nevertheless, online information services do currently provide some targeting capability. As of August 1996, both Yahoo! [24] and Excite [9] offered advertisers three options: general rotation, geographic or content targeting, and keyword-based targeting. With "general rotation," banner advertisements rotate randomly through user searches and browsing on the site. The Honda advertisement that appeared on the Organic Gardening page in Appendix 3 was probably in general rotation. Restricted rotations permit advertisers to purchase space in specified content areas or by geographic region. For example, financial institutions can limit the exposure of their banner advertisements to users searching or browsing Business categories, and Canadian advertisers can choose to have their banner advertisements shown only to users who are searching or browsing in the Yahoo! Canada site. These two options are analogous to the targeting offered by traditional media such as newspapers, magazines, television, and radio [4].

The third option, keyword-based targeting, makes greater use of the targeting potential of information services. A company can buy keywords so that whenever a user enters one of those keywords during a search, s/he will be exposed to the company's banner advertisement. This ensures that the banner advertisement is presented only to people with a

demonstrated interest in the area. For instance, a marketer of golf equipment might buy the keyword "golf." Every time a user enters "golf" in a search, a banner advertisement for the equipment would appear. This is analogous to the more precise targeting provided by magazines.

While these are useful strategies, they fail to take full advantage of the targeting potential of banner advertising. Current technology provides the capability to develop sophisticated and detailed profiles of individual users of information services based on individual characteristics and past patterns of behavior in using the information service. The next section proposes and describes informally two versions of a model for targeting banner advertising by using the information technology on which an online information service is built.

3. A Model for Targeted Advertising

In traditional media, the quality of the information available constrains an advertiser's ability to target advertising effectively and efficiently. For example, many media buying decisions are based on data provided by research bureaus such as the Audit Bureau of Circulations (ABC), Business Publication Audit of Circulation (BPA), Arbitron, and A.C. Nielsen, which collect data on the demographics and media habits of consumers, and sometimes on product usage and brands [4]. These survey data are cross-tabulated to develop a profile of the audience of each media vehicle. The audience profile is then compared to the target audience profile identified by the advertiser to determine where there is a good match. For instance, an automobile manufacturer might identify the target audience for an advertisement for a particular model of car as middle-income females, 18 to 34, with busy lifestyles. Based on research bureau data, as well as the experience and judgement of the media planner, media vehicles with good reach in that demographic group would be chosen. Realistically, though, this type of targeting is usually very approximate. For instance, no matter how well the media vehicle audience profile matches the target audience profile, it is likely that only a portion of the audience would be in the market for a new car.

Online banner advertising may be able to overcome this problem. It is possible to target users very precisely because data can remain associated with individuals, so advertisers can select exactly the users to whom they wish their advertising to be exposed. It may be possible, for example, to identify which users will be in the market for a new car in a particular year. The remainder of this section describes two versions of a model for targeting banner advertising by taking advantage of the technological capabilities of the online environment. The model is designed to be appropriate for use by information services which sell

advertising space.

3.1. Basic version

The basic version of the model requires that users be assigned unique identifiers (e.g., user accounts) when they first connect to the information service. Subsequently, they provide these identifiers each time they connect. Users also complete an online questionnaire the first time they use the information service. (Incentives to complete the questionnaire may be provided by informing users that the information will be used to filter out advertising for products in which they are likely not to be interested.) The questionnaire allows data to be collected on several dimensions, including: (1) demographic attributes such as geographic location, income, family lifecycle stage, occupation, and sex; (2) psychographic attributes such as travel patterns and hobbies; and (3) product and brand usage attributes. This element of the basic model permits a banner advertisement to be directed to users (and only those users) who fit certain criteria, assuming data were collected on relevant attributes. For instance, a banner advertisement for baby strollers could reach parents of children under five years old—and only individuals in that group.

In contrast, research bureau data uses demographic correlates (e.g., males and females, 18 to 34) to identify media vehicles that attract a relatively large proportion of the people in the identified demographic group [4]. The media vehicles thus chosen may miss members of the target group (e.g., older parents) and reach consumers not in the target group (e.g., people who are between 18 and 34 but do not have young children). Even audience data based on cross-tabulations, while they supply information on more variables, still cannot isolate individuals who are in the target audience. (For example, research bureau data may allow an advertiser to identify a magazine whose audience includes a large number of people between 18 and 34 who have young children, but there will still be some readers who are not in the target market.)

The second element of the basic model involves eliciting the target audience profile from advertisers. An advertiser can specify a target audience using any number of attributes about which data have been collected. These can be expressed conjunctively and/or disjunctively. For example, a specification may indicate that an advertisement is to be presented to all users who (1) have household incomes over \$50,000, and (2) either work in a job that involves travel at least four times per year or have travelled on vacation in at least four of the past five years.

In this version of our model, the questionnaire determines the data collected about each user. The content of the questionnaire will vary depending on the nature of the information service, expected users, and expected

advertisers. However, it is imperative to design the instrument carefully, in consultation with advertisers based on anticipated relevant target audience attributes.

The final element of the model consists of a mechanism to select banner advertisements to display to users. The target audience profiles supplied by advertisers provide a screening mechanism over users. Each time a user connects, his/her profile is compared to all target audience profiles from all advertisers. The user's profile will actually match some subset of those profiles. If the number of matches is small (and the session is long), it will be feasible to display all banner advertisements associated with the matched profiles during the user's session. However, if the number of matches is larger (or the session is short), presenting all advertisements associated with the matched profiles may overwhelm the user. In such a case, it will be necessary to present only a selection of the identified target advertisements. A rationing system would be needed so that users are not deluged with banner advertisements while advertisers are assured of access to users who match the target audience profile.

In summary, the basic model has three elements: individual user profiles, individual advertisement target audience profiles, and a selection mechanism for presenting advertisements to specific users who match the target audience profile. This framework potentially eliminates wasted exposures and provides the capability to reach every single user who matches the target audience profile (this may not be realized if a rationing system is used). Users also benefit, since they will see advertisements only for products likely to be of interest to them.

3.2. Enhanced Version

The basic version of the model relies on users completing a questionnaire when they initially use an information service. This is a straightforward mechanism to collect data about user characteristics for the purpose of targeting advertisements. A similar approach has been incorporated in a commercial product for use with online catalogs to direct shoppers to products in which they are interested [15]. However, the advantage of simplicity is offset by several potential limitations. First, such information may become outdated sometimes quickly, as user preferences and characteristics change. To some extent, information can be kept up-to-date by either readministering the questionnaire periodically or giving the user the opportunity to update her/his information (e.g., by a menu option or hypertext link) each time s/he connects to the information service. However, each of these strategies is intrusive and may impose an unwarranted burden on users in order to maintain currency of information.

A second, and perhaps more serious limitation of the

questionnaire strategy is that it is subject to two potential types of bias. First, the questionnaire designer will want to identify as many user attributes relevant to potential advertisers as possible. As the number of attributes increases, so does the length of the questionnaire, creating the possibility of higher mortality in completing the questionnaire (especially since it may be more difficult to induce users to complete it because they are both physically and psychologically remote), thereby increasing the potential nonresponse bias [7]. Second, the questionnaire method is plagued with well-known problems, such as errors due to inaccurate recall, telescoping, social desirability concerns, and cognitive biases, as well as ambiguity, intimidation, confusion, and incomprehensibility [2].

In view of these potential problems, it is appropriate to enhance the model so that it does not rely on user self-reports, can accommodate changing user characteristics and preferences, and is less constrained by the choice of questions. Fortunately, information technology may provide assistance in each of these areas.

Current technology allows a considerable amount of data about user search activities (both deliberate search and browsing) to be collected unobtrusively and analyzed to determine patterns. (We are dealing here only with the capabilities of the technology, not with the ethical issues such capabilities raise. However, we recognize that ethical issues must be considered explicitly in the design of systems based on our model. For instance, we believe users should be aware that such information may be collected, and how it may be used, and consent to this activity before using an information service.) In the enhanced model, we propose that patterns of search and browsing behavior exhibited by users while using an information service determine which advertisements are shown to that user during current or future sessions. In the remainder of this section, we provide a general overview of this approach.

As before, this model relies on assigning a unique identifier to each user for recording her/his searching and browsing activities while using the information service. Each session constitutes a "record", consisting of data such as: sites visited in order; pattern of navigation through a hierarchical category structure (as in Yahoo!); choice of search terms in keyword-based searches; and reaction to previously exposed targeted banner advertisements (e.g., which linked Web sites are selected and visited by the user and which ones ignored). The aggregate of such records for each user provides a profile from which preferences can be implicitly generated. As a simple example, if a user has made several searches using keywords such as "Atlantic salmon" and "fly fishing", and has visited the site of the *Angling Club Lax-a-of-Iceland* (http://www.ismennt.is/fyr_stofn/lax-a/uk/angl_uk.html), s/he may be targeted for a banner advertisement for a fishing

lodge in Alaska. However, if a user has previously been exposed to the same or similar banner advertisements but has not visited linked Web sites when there was an opportunity to do so, s/he may not be shown these banner advertisements in future.

This version of the model has the advantage of transparency. A user simply visits a service for whatever purpose s/he has in mind. Data are collected unobtrusively in the course of the visit. Moreover, the data reflect actual user behavior, rather than attitudes, intentions, or reported behavior captured through a questionnaire. Hence, the quality of data derived from user behavior should be superior to that of questionnaire data, for purposes of targeting advertisements.

A disadvantage of this model is the preparatory work involved on two fronts. First, it is not clear how to structure the data collected during visits so that useful information can easily be coded for storage and later extraction. Research is needed to develop useful and efficient coding mechanisms for storing such data as sequences of visits and search terms used. We expect this can be handled using conventional database structures such as relations (tables); however, the design of a relational database for this purpose is itself a distinct research issue. Second, the ability to store the required data does not necessarily mean useful information can be extracted from it. Further research is required to determine the types of analyses that yield insights into user characteristics and preferences hidden in the data.

The enhanced model should be used in conjunction with the basic model. A questionnaire may be very effective for identifying various demographic data relevant to advertisers but impossible to ascertain simply from users' online search and browsing behavior. However, since demographic data has limitations for effectively targeting consumers of most products, the enhanced model of data collection may yield complementary data on preferences from patterns of online search and browsing behavior.

The next section describes an implementation architecture for the basic version of the model. Extensions that support the enhanced version of the model remain as future research.

4. An Implementation Architecture

The architecture required to implement the basic version of the model consists of two parts: data structure to represent user profiles and target audience profiles, and an algorithm to select banner advertisements to display to a user. This section describes these components.

4.1. Data Structure

To target banner advertisements, two types of profiles are needed: profiles describing users of the information service; and profiles describing the target audience for advertisements, as defined by advertisers. Each profile can be modeled as a set of attributes.

We assume there is a finite "universe" of attributes, $A = \langle a_1, \dots, a_N \rangle$, that may potentially characterize users or target audience members.

4.1.1. User Profile. Each user, u_i , of the service can be described by a record consisting of values of the universe of attributes, $R_i = \langle a_1(u_i), \dots, a_N(u_i) \rangle$, where $a_n(u_i)$ ($n=1, \dots, N$) denotes the value of attribute a_n for user u_i . This may be implemented in a relational database in which a table is defined whose primary key is a user identifier, and remaining attributes are those in A . Each row in the table contains the profile of one user. (A more elaborate data structure is needed to support the enhanced model, since data must also be kept about the pattern of behavior of a user over one or more sessions.) All attributes need not be applicable or relevant to a particular user; hence, null values are permitted.

A simple example serves to illustrate this structure. Consider a universe consisting of three attributes: age, income, and number of dependents. Suppose there are two users of a service. When those users have completed a profile questionnaire, the resulting data may be stored in a relational table as:

USER			
user_id	age	income	dependents
u ₁	26	34000	0
u ₂	45	54000	2

4.1.2. Target Audience Profile. A target audience profile is associated with each banner advertisement. A profile may be expressed as:

- (1) *A characterization of an "ideal" target audience member.*

Such an ideal can be described by a record consisting of values of the universe of attributes, $T_i = \langle t_1, \dots, t_N \rangle$, where t_n ($n=1, \dots, N$) is a specific value of attribute a_n . Some values may be null, indicating that any values of those attributes are permitted for the ideal; and/or a

- (2) *A characterization of the "acceptable" target audience.* Generally, an advertiser is interested in reaching those within specified ranges of the attributes of interest. Given N attributes of interest, acceptability can be thought of as a region in N -dimensional space. This region can be defined by specifying ranges of acceptable values for various attributes in the universe.

Combinations of attributes may be expressed conjunctively, indicating that users in the target region must satisfy all the conditions or restrictions; disjunctively, indicating that acceptable users must satisfy one of a set of conditions; or using a combination of disjunctions and conjunctions.

Note that "distance" from the ideal point may become relevant if an advertiser has to choose a subset of users whose profiles fall within the acceptable region.

Operationally, profiles for ideal or acceptable users can be maintained in a relational database structure. In the case of ideal profiles, a table can be defined in which each row describes the ideal target audience member for each advertisement. The primary key for this table consists of an identifier for the advertisement, while the remaining attributes are those of the universe of attributes of interest. Since not all attributes may be relevant in specifying an ideal, null values are permitted.

To illustrate, consider a simple example in which there are two advertisements, each with a different target audience profile, designated T_1 and T_2 . The ideal target profile for T_1 is users aged 35 with incomes of \$50,000 (no restrictions on number of dependents), while that for T_2 is users aged 25 with incomes of \$25,000 and no dependents. These profiles are shown in the following relational table.

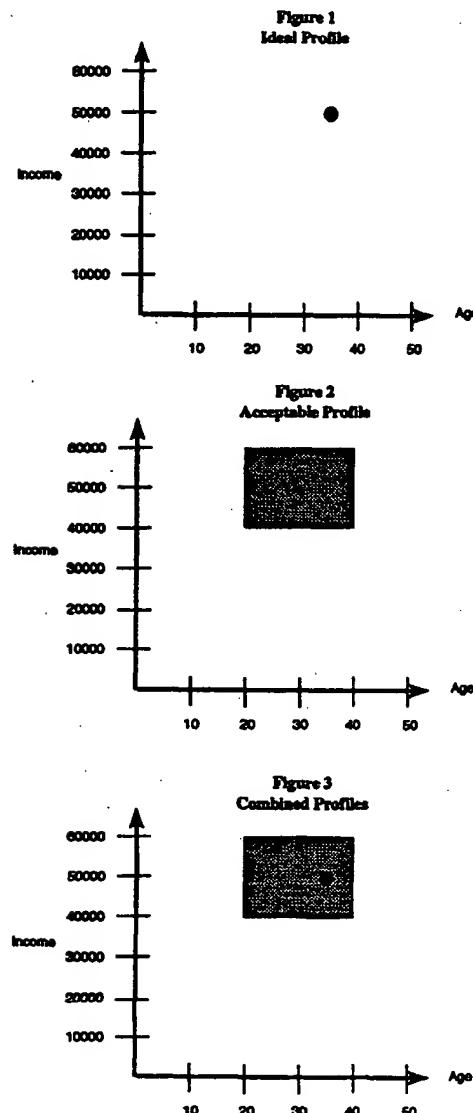
TARGET			
ad_id	age	income	dependents
T ₁	35	50000	0
T ₂	25	25000	0

In this example, the space of profiles is three-dimensional. Since the "dependents" attribute is not relevant in describing the target audience of the first advertisement, the ideal profile can be depicted as a point in two-dimensional space, as shown in Figure 1. In the case of target audience profiles based on attribute values within a range, the necessary data structure can be provided by a table whose rows describe the acceptable ranges of specified attributes for each profile. Each row in this table provides a lower and an upper bound on a specified attribute for a specified profile. The primary key of this table consists of the identifier of the profile plus the name of the attribute. To illustrate, consider a variation of the example above. Suppose the profile T_1 is no longer age = 35, income = \$50,000, but is relaxed: the advertiser will accept any user between 20 and 50 with an income between \$40,000 and \$60,000. Similarly, suppose the profile T_2 is relaxed to encompass ages from 20 to 30 and incomes from \$20,000 to \$30,000. Such profiles might be stored as follows:

RANGE	ad_id	attribut	lower	upper
T ₁	age	20	50	
T ₁	income	40000	60000	
T ₂	age	20	30	
T ₂	income	20000	30000	

In this case, the acceptable profile can be depicted as a region in two-dimensional space. Figure 2 shows the profile for T₁.

It is possible that both ideal and acceptable profiles could be generated by the same advertiser. By overlaying Figures 1 and 2, shown in Figure 3, we note that the ideal point need not lie at the geometric center of the acceptable region.



To handle measures of "distance" from an ideal, ranges of values on relevant attributes can be replaced with advertiser-specified information about the acceptable distributions of values over attributes. For instance, an advertiser may specify a mean (ideal) and standard deviation for an attribute if "acceptability" is normally distributed about a central value. Other measures of central tendency and dispersion may be appropriate for attributes in which the range of acceptability is quantified differently. The data structure of the RANGE table can be modified to accommodate this additional complexity.

4.2. Selecting Advertisements for Users

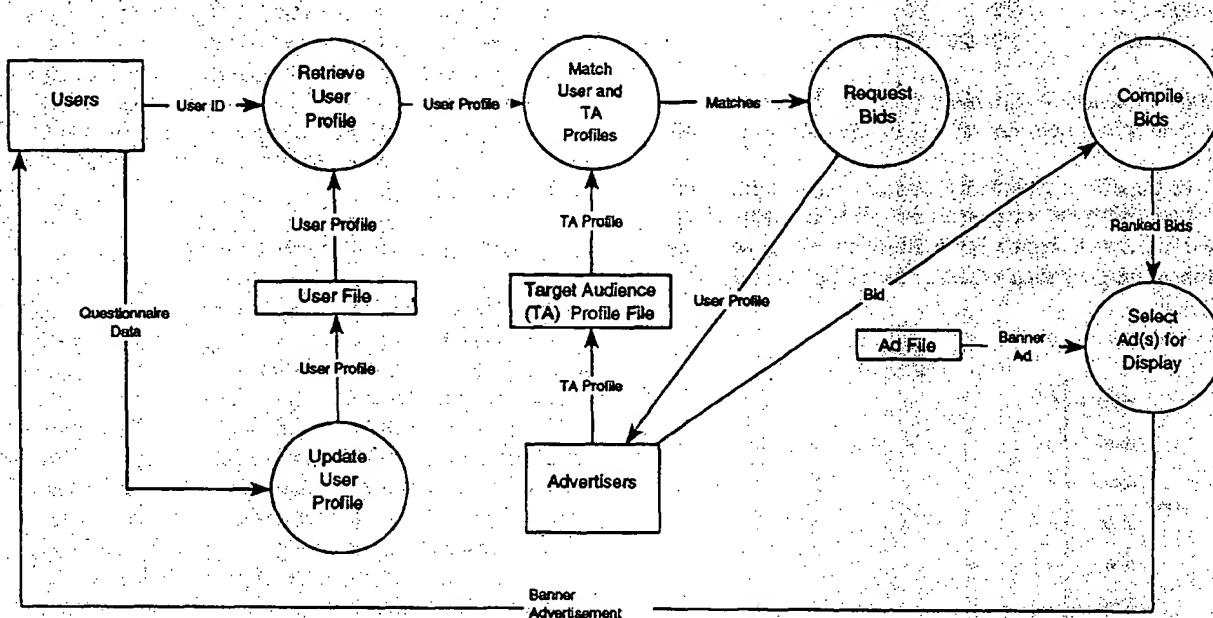
The primary challenge in effectively and efficiently targeting banner advertising is matching user profiles with target audience profiles. Figure 4 uses a data flow diagram to depict the matching process described below.

When a registered user visits an information service, his/her profile is retrieved. This profile is then compared with the target audience profiles of banner advertisements currently being run by the information service. Each target audience profile is associated with a banner advertisement. For each target audience profile, if there is no match with the user profile, the associated advertisement is dropped from further consideration.

After the comparisons are completed, a set of matched target audience profiles from a variety of advertisers remains. If this set is small, it may be feasible to show all the associated banner advertisements to the user during the session. In general, though, it will be necessary to select some subset of advertisements from the matched set to display to the user. We envision that the advertisers whose advertisements are in the matched set will compete for the opportunity to have their banner advertisements displayed to the user.

The concept of acceptable regions in target audience profiles provides a basis for competition. Profiles accommodate the possibility that some users within the region of acceptability may be more desirable to an advertiser than others. Hence, a distance metric capturing the relative desirability of a user with respect to an ideal profile is possible. It is not the purpose of this paper to propose or evaluate metrics. However, recognizing a notion of distance allows the possibility for advertisers to "bid" for the opportunity to display an advertisement to a user. Such bids would be determined by the advertiser, based on variables such as the user profile (to determine the distance from the ideal target audience profile) and advertising budget. It may be feasible to automate this by having software agents associated with each advertisement that would calculate the distance measure for the user and formulate a bid based on this, in addition to other

Figure 4
Selecting Advertisements for Users



information such as whether the user had seen this advertisement, or other advertisements for the same or similar products, in previous sessions (information which could be carried as part of the user profile).

When bids are received, they can be ranked. The banner advertisement corresponding to the winning bid is displayed to the user. Other advertisements may be displayed according to their ranking if there is an opportunity to display additional advertisements (e.g., if the user engages in several search or browse activities during a session).

This architecture provides guidance for implementing the basic version of the model. We present next a simple example showing how the architecture operates.

4.3. Example

Consider the relational database tables USER, TARGET, and RANGE presented earlier. Suppose first that the user with profile u_1 connects to the information service. This user's profile, consisting of the database record $\langle u_1, 26, 34000, 0 \rangle$ is retrieved from the USER table. Next, target audience profiles T_1 and T_2 are retrieved from the TARGET table. These identifiers determine the attributes whose profile ranges have to be selected from the RANGE table. Next, the age range for T_1 , namely $(20, 30)$, is retrieved from RANGE. Since the age value of u_1 is 26, there is a match on

this criterion. So, the salary range for T_1 , $(40000, 60000)$ is retrieved. Since the user u_1 does not match this criterion of the target audience profile (salary is 34000), the advertisement corresponding to the profile T_1 will not be shown to u_1 . Applying the same operations to the target audience profile T_2 , u_1 would not be exposed to the advertisement corresponding to T_2 since the income of u_1 (34000) is greater than the upper bound of 30000 specified in the target audience profile T_2 . Thus the user with profile u_1 would not be exposed to any banner advertisements when s/he used the information service. This is efficient, since showing either banner advertisement to the user with profile u_1 would entail a cost and constitute a wasted exposure.

Suppose now that the user with profile u_2 connects to the information service. This user's profile, $\langle u_2, 45, 54000, 2 \rangle$, is first retrieved from the user profile file. The target audience profiles T_1 and T_2 are then retrieved. Applying the matching algorithm, a match will be found between u_2 and the profile T_1 . (Note that the target audience profiles in our example do not specify restrictions on number of dependents; hence, any values are permitted on this attribute.) However, there is no match between u_2 and T_2 since the income of u_2 (54000) is beyond the upper bound of income for T_2 (30000). Hence, the user with profile u_2 will be exposed to the advertisement corresponding to the target audience profile T_2 .

This example does not show the full scope of the model, since there is no case where there are two or more target audience profiles that match a particular user profile. To illustrate this, consider an additional target audience profile having only the condition that user income must be at least 50000. This profile would require adding the following record to the RANGE table: <P_i, income, 50000, ->. Now if the user with profile u_j connects to the information service, a match with both T_i and T_j will be found. In this case, the advertisers (or software agents) responsible for T_i and T_j will be contacted and provided with the profile u_j. Each advertiser (agent) will prepare a bid indicating how much it is willing to pay to have the banner advertisement corresponding to its profile exposed to the user with profile u_j. These bids are compiled and returned to the information service, where they are ranked. If we allow that the user will be shown only one advertisement, the one which placed the highest bid will be chosen.

In summary, this model makes use of rich, multiattribute data at the individual user level in determining whether each one will be exposed to a particular banner advertisement. This leads to more effective and efficient targeting than is possible using strategies such as general rotation, which does not use data at the individual level, and restricted rotation or keyword search, which rely on only a single data item about an individual in determining which banner advertisement(s) to present to the user. It is also more effective and efficient than targeting in the traditional media, which does not use any data at the individual level.

5. Future Research

This paper has presented a framework for leveraging information technology to target online banner advertising more effectively to benefit both users (who would be exposed only to advertising that is very probably of interest to them) and advertisers (whose advertisements would reach only those users who fit the target audience profile). This framework is, however, merely a starting point. Additional research on several fronts is needed before its potential can be realized. Several specific research concerns have already been noted. In addition, there are more general issues.

First, a system supporting the basic version of the model, based on the implementation architecture presented in this paper, should be implemented.

Second, an implementation supporting the enhanced version of the model is needed. This will require research to develop a more sophisticated database structure that can preserve users' searching and browsing behavior over time. In addition, techniques for detecting patterns of behavior are needed.

Third, both theoretical and empirical research is needed to explore agent bidding in the context of the framework

proposed in this paper.

Fourth, empirical work needs to be done to evaluate the relative effectiveness and efficiency of this framework. A priority should be to compare the (1) the basic version of the model, (2) the enhanced version of the model, and (3) existing approaches to targeting advertisements. For instance, it would be interesting to test whether placing a banner advertisement on a relevant page (e.g., an advertisement for a new movie on the Entertainment page of an information service) would be more or less effective than directing the same advertisement to individual users selected on the basis of their answers to a questionnaire (i.e., the simple version of the model) or their search and browsing behavior (i.e., the enhanced version of the model).

Finally, the utility of this framework in other online contexts should be investigated. For instance, this approach could be used in developing Web sites that are more useful to visitors. Visitors with different profiles could automatically be shown different pages more likely to be of interest to them, eliminating the need for them to search the Web site for the information they desire.

6. Conclusion

Cyberspace is a new medium for advertising. In 1994, Edwin Artzt, chairman of Procter & Gamble, the largest advertiser in the United States, warned advertising agencies to "get their interactive act together" [6, p. 75]. As the advertisements in Appendices 1 and 3 show, even major advertisers and their agencies may not be taking full advantage of the opportunity to target their online banner advertising. The information technology that makes the WWW possible also allows the unobtrusive collection of detailed information about user interests based on their online searching and browsing. Advertisers should not assume that the same constraints that make media planning in traditional media a very inexact science also apply to online advertising.

7. References

1. Armstrong, Arthur and John Hagel III (1996), "The Real Value of On-line Communities," *Harvard Business Review* 74(3), 134-141.
2. Barnes, James G. (1991), *Research for Marketing Decision Making*. Toronto, ON: McGraw-Hill Ryerson.
3. Berthon, Pierre, Leyland F. Pitt, and Richard T. Watson (1996), "The World Wide Web as an Advertising Medium: Toward an Understanding of Conversion Efficiency," *Journal of Advertising Research* 36(1), 43-54.
4. Bovée, Courtland L., John V. Thill, George P. Dovel, and Marian Burke Wood (1995), *Advertising Excellence*. New York, NY: McGraw-Hill.
5. *Business Week*, June 3, 1996, 44.
6. Castro, Janice (1995), "Just Click to Buy," *Time*, Spring, 74-

- 75.
7. Churchill, Gilbert A. (1991), *Marketing Research: Methodological Foundations*, Fifth Edition. Orlando, FL: Dryden.
 8. Cortese, Amy (1995), "Cyberspace: Crafting Software that Will Let you Build a Business Out there," *Business Week* (February 27), 78-86.
 9. Excite (1996), *Excite Advertising Rate Card*. [<http://www.excite.com/info/advertising.html>]
 10. Greenwald, John (1995), "Battle for Remote Control" *TIME* Special Issue: Welcome to Cyberspace (Spring), 71-73.
 11. Gupta, Sunil and Jim Pitkow (1995a), "Consumer Survey of WWW Users: Preliminary Results from the Third Survey," University of Michigan Business School. [<http://www.umich.edu/~sgupta/hermes>]
 12. Gupta, Sunil and Jim Pitkow (1995b), "Consumer Survey of WWW Users: Preliminary Results from the Fourth Survey," University of Michigan Business School. [<http://www.umich.edu/~sgupta/hermes>]
 13. Hoffman, Donna L. and Thomas P. Novak (1996), "Marketing in Hypermedia Computer-Mediated Environments: Conceptual Foundations," *Journal of Marketing* 60(3), 50-68.
 14. Hoffman, Donna L., Thomas P. Novak, and Patrali Chatterjee (1995), "Commercial Scenarios for the Web: Opportunities and Challenges," *Journal of Computer-Mediated Communication, Special Issue on Electronic Commerce* 1(December). [<http://shum.huji.ac.il/jcmc/vol1/issue3/hoffman.html>]
 15. Information Week, "One-to-One Marketing Made Easy", *Information Week*, December 18, 1995, 110.
 16. Kahn, Robert E. (1994), "The Role of Government in the Evolution of the Internet," *Communications of the ACM* 37(8), 15-19.
 17. Karlgaard, Rich (1995), "Bozo Explosion," *Forbes ASAP* (December 4), 9.
 18. Landau, Susan, Stephen Kent, Clint Brooks, Scott Charney, Dorothy Denning, Whitfield Diffie, Anthony Lauck, Douglas Miller, Peter Neumann, and David Sobel (1994), "Crypto Policy Perspectives," *Communications of the ACM* 37(8), 115-121.
 19. Leiner, Barry M. (1994), "Internet Technology," *Communications of the ACM* 37(8), 32.
 20. Ottawa Citizen, April 29, 1996, A10.
 21. Press, Larry (1994), "Commercialization of the Internet," *Communications of the ACM* 37(11), 17-21.
 22. Quelch, John A. and Lisa R. Klein (1996), "The Internet and International Marketing," *Sloan Management Review* 37(3), 60-75.
 23. Tetzeli, Rick (1994), "The Internet and Your Business," *Fortune* 129 (March 7), 86-88+.
 24. Yahoo! (1996), *Sponsorship Program Opportunities* [<http://www.yahoo.com/docs/pr/adop.html>]

8. Appendices

Appendix 1



Appendix 2



Appendix 3



**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER: _____**

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.